

An Approach to Reduce Communication for Multi-agent Mapping Applications

Michael E. Kepler¹ and Daniel J. Stilwell²

Abstract—In the context of a multi-agent system that uses a Gaussian process to estimate a spatial field of interest, we propose an approach that enables an agent to reduce the amount of data it shares with other agents. The main idea of the strategy is to rigorously assign a novelty metric to each measurement as it is collected, and only measurements that are sufficiently novel are communicated. We consider the ideal scenario where an agent can instantly share novel measurements, and we also consider the more practical scenario in which communication suffers from low bandwidth and is range-limited. For this scenario, an agent can only broadcast an informative subset of the novel measurements when the agent encounters other agents. We explore three different informative criteria for subset selection, namely entropy, mutual information, and a new criterion that reflects the value of a measurement. We apply our approach to three real-world datasets relevant to robotic mapping. The empirical findings show that an agent can reduce the amount of communicated measurements by two orders of magnitude and that the new criterion for subset selection yields superior predictive performance relative to entropy and mutual information.

I. INTRODUCTION

Gaussian process regression has been employed to model various nonlinear spatial phenomena, such as terrain elevation [1], temperature [2], and radio signal strength [3], as it combines past measurements and sensor uncertainty to generate a probabilistic model that given a set of query inputs provides predictions with associated uncertainty. The predictive model has a closed form so sampling methods, which can be computationally expensive, are not required to approximate the predictive model. Additionally, Gaussian process regression admits a continuous-domain representation, meaning that query points can reside anywhere in the input domain and need not be restricted to a fixed resolution.

Gaussian process regression has been used as the map representation in collaborative mapping applications for numerous multi-agent systems, of which a characteristic subset includes [4]–[8]. Multi-agent systems subject to limited communication range has been addressed by various authors [9]–[11], among others. However, few studies address multi-agent systems that must cope with low communication bandwidth. In this paper, we propose an approach that translates

the coordination approaches like those found in [4]–[8] to the setting in which communication bandwidth is limited.

The coordination of multi-agent systems [4]–[8], among others, aims to compute paths for each agent so as to optimize an objective function, such as conditional entropy [8], mutual information [6], [9], and variance [5], [7], all of which rely on a representation of the map. Thus, if agents are not aware of the measurements that other agents have collected and can not update their map, it is possible that planning paths for agents with an obsolete map representation could lead to arbitrarily poor mission performance.

In [4]–[8] the ability of an agent to share its measurements is of no concern as agents can communicate data as frequently as desired, provided the agent is within communication range of the other agents. However, applications in environments characterized by low bandwidth can prohibit an agent from sharing data as frequently as desired. In subsea applications, for example, the communication channel drastically attenuates electromagnetic signals so submerged agents must resort to low frequency acoustic communication [12]. In order to translate the coordination approaches that assume high bandwidth, like those in [4]–[8], to low bandwidth environments we propose that an agent share only those measurements which have a significant impact on the map in some well-defined sense. While our work is motivated by multi-agent mapping, coordination strategies are not addressed specifically in this paper. Instead, we propose a rigorous approach for an agent to reduce the amount of data that it shares with others so as to facilitate collaborative mapping in communication-limited environments.

Our approach uses the sparse online Gaussian process regression algorithm [13] as it assigns a rigorous novelty metric to each measurement and only sufficiently novel measurements are retained for regression. This novelty metric is used to decide which measurements should be shared. We consider both the case that communication is not range limited and the case that it is limited. For the latter setting, we assume the collection of novel measurements is too large to be transmitted in its entirety and an agent must resort to sharing an informative subset of the novel measurements. For subset selection criteria we consider the widely-used information-theoretic measures entropy and mutual information [2], and introduce a new criterion that reflects the value of the measurement. To demonstrate the efficacy of our approach, we apply it to a collection of three real-world data sets: a bathymetric dataset, magnetic field dataset, and a radio signal strength dataset.

*This work was supported by the National Defense Science and Engineering Graduate (NDSEG) fellowship program and the Office of Naval Research via grants N00014-18-1-2627 and N00014-19-1-2194

¹Michael E. Kepler is an NDSEG fellow and a Ph.D. candidate in the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA 24060, USA mkepler@vt.edu

²Daniel J. Stilwell is a Professor in the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA 24060, USA stilwell@vt.edu

II. RELATED WORK

Our work on reducing communication among agents is closely related to the "Single Agent in a Team Decision" problem [14], where an individual agent gathers new information and must decide whether or not to share the information with the other agents. In principle, this requires the agent to quantify the utility of communicating the information. To do so, numerous researchers [14]–[18] cast the problem as a Markov decision process and augment the action space of an agent with the ability to share new information. To decide whether or not to share new information, the agent computes the expected improvement in reward due to communication, and if the improvement is greater than the cost of communication, then the agent shares the new information with other members of the team. Moreover, the works [15], [18] consider the case in which bandwidth is limited so that an agent can only transmit k new observations. The goal then is to find the k observations that most increase the expected team reward if communicated. The key limitation of the Markov decision process framework is that computing the expected reward is computationally prohibitive [16], and hence approximation techniques are often required to compute the expected utility of communication.

Our approach to reduce communication, on the other hand, does not rely on the Markov decision process framework to compute the utility of a measurement. We maintain the underlying principle that an agent assigns utility to measurements, and only measurements of sufficient utility are communicated. However, we consider only the estimation problem that arises from multiple agents approximating a spatial field using Gaussian process regression. We use the novelty of a measurement as a proxy for communication utility, and leverage the closed-form expressions given by the sparse online Gaussian process regression framework [13]. Although our approach is computationally efficient, we do not consider a rich set of decision processes that can be modelled in an Markov decision process setting. To the best of the authors' knowledge, this is the first work that addresses the reduction of communication for multi-agent systems that use a Gaussian process to estimate a spatial field, such as those in [4]–[8].

III. GAUSSIAN PROCESS REGRESSION

Let $S_t = \{(\mathbf{x}_i, z_i)\}_{i=1}^t$ denote the set of t input-measurement pairs used for regression. In the standard Gaussian process regression setting [19, Ch. 2.2], each measurement $z_i \in \mathbb{R}$ is related to its associated input $\mathbf{x}_i \in \mathcal{X}$ via $z_i = f(\mathbf{x}_i) + \varepsilon_i$ for $i = 1, \dots, t$. The additive noise ε_i is assumed to be independent and identically distributed so that $\boldsymbol{\varepsilon} = [\varepsilon_1, \dots, \varepsilon_t]^\top$ arises from a multivariate normal distribution with mean vector $\mathbf{0}_t$ and covariance matrix $\sigma_n^2 \mathbf{I}_t$, where \mathbf{I}_t is the t -dimensional identity matrix. Given a location of interest \mathbf{x}_* the objective is to predict the underlying spatial field at that input $f_* \triangleq f(\mathbf{x}_*)$. Conditioned on the collection of measurements $\mathbf{z}_t = [z_1, \dots, z_t]^\top$ and associated locations $\mathbf{X}_t = \{\mathbf{x}_1, \dots, \mathbf{x}_t\}$, the posterior predictive distribution for f_*

is Gaussian, specified by the mean $\mu_{\mathbf{x}_*|S_t}$ and variance $\sigma_{\mathbf{x}_*|S_t}^2$

$$\mu_{\mathbf{x}_*|S_t} = \mathbf{k}(\mathbf{x}_*)^\top (\mathbf{K}_t + \sigma_n^2 \mathbf{I}_t)^{-1} \mathbf{z}_t, \quad (1a)$$

$$\sigma_{\mathbf{x}_*|S_t}^2 = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{x}_*)^\top (\mathbf{K}_t + \sigma_n^2 \mathbf{I}_t)^{-1} \mathbf{k}(\mathbf{x}_*), \quad (1b)$$

where the i^{th} element of the vector $\mathbf{k}_t(\mathbf{x}_*) \in \mathbb{R}^t$ is $[\mathbf{k}_t(\mathbf{x}_*)]_i = k(\mathbf{x}_i, \mathbf{x}_*)$, and the i, j^{th} element of the gram matrix $\mathbf{K}_t \in \mathbb{R}^{t \times t}$ is $[\mathbf{K}_t]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$.

The covariance function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is user-specified, and given \mathbf{x} and \mathbf{x}' , $k(\mathbf{x}, \mathbf{x}')$ encodes the user's belief about the covariance between the measurements z and z' prior to collecting any measurements. Throughout this paper, we use the squared exponential covariance function

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left\{ -\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2} \right\}, \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean norm, σ_f^2 denotes the signal variance, and ℓ denotes the length-scale parameter.

There are various ways to select the hyperparameters $\boldsymbol{\theta} = \{\ell, \sigma_f, \sigma_n\}$. One approach is to select the set of hyperparameters that maximize the log likelihood

$$\begin{aligned} \ln p(\mathbf{z}_t | \mathbf{X}_t, \boldsymbol{\theta}) &= -\frac{1}{2} \mathbf{z}_t^\top (\mathbf{K}_t + \sigma_n^2 \mathbf{I}_t)^{-1} \mathbf{z}_t \\ &\quad - \frac{1}{2} \ln |\mathbf{K}_t + \sigma_n^2 \mathbf{I}_t| - \frac{t}{2} \ln 2\pi, \end{aligned} \quad (3)$$

where $|\mathbf{A}|$ denotes the determinant of the matrix \mathbf{A} . The set of optimal hyperparameters are obtained by computing the partial derivatives of (3) with respect to $\boldsymbol{\theta}$ and using a gradient-based optimizer. For further details see [19, Ch. 5.4]. Alternatively, if knowledge from an initial set of data or expert knowledge is available, then the hyperparameters can be tuned manually. For the empirical studies in this paper, we use the hyperparameters that maximize the log likelihood, listed in Table I.

For applications with a large (e.g. $t > 10,000$) number of measurements, prediction using the standard Gaussian process regression equations (1) becomes problematic as the standard Gaussian process algorithm is a batch algorithm, meaning that whenever a new measurement z_{t+1} is collected, the old matrix inverse $(\mathbf{K}_t + \sigma_n^2 \mathbf{I}_t)^{-1}$ is discarded from memory and the new inverse $(\mathbf{K}_{t+1} + \sigma_n^2 \mathbf{I}_{t+1})^{-1}$ is computed and stored in memory. Moreover, the computational cost of the matrix inverse $(\mathbf{K}_t + \sigma_n^2 \mathbf{I}_t)^{-1}$ typically scales as $\mathcal{O}(t^3)$, prohibiting the use of the standard Gaussian process prediction equations for real-time applications. To circumvent these computational costs we use the sparse online Gaussian process regression algorithm [13] which alleviates the computational cost associated with standard Gaussian process regression.

In the remainder of this section we present the sparse online Gaussian process regression algorithm and the derivation of the novelty metric as they play an essential role in our approach, but for a more detailed exposition we refer the reader to [13].

The sparse online Gaussian process regression algorithm is well-suited for real-time mapping applications due to the

fact that (i) the predictive mean and variance are expressed in terms of the recursively updated variables $\boldsymbol{\alpha}_t \in \mathbb{R}^t$ and $\mathbf{C}_t \in \mathbb{R}^{t \times t}$

$$\boldsymbol{\mu}_{\mathbf{x}_*|S_t} = \mathbf{k}_t(\mathbf{x}_*)^\top \boldsymbol{\alpha}_t, \quad (4a)$$

$$\boldsymbol{\sigma}_{\mathbf{x}_*|S_t}^2 = k(\mathbf{x}_*, \mathbf{x}_*) + \mathbf{k}_t(\mathbf{x}_*)^\top \mathbf{C}_t \mathbf{k}_t(\mathbf{x}_*), \quad (4b)$$

and (ii) the number of measurements explicitly maintained for regression is limited by retaining in memory only those data points that are sufficiently novel, where the notion of novelty is made precise in terms of projection error. Conceptually, inputs are first mapped to a linear Hilbert space of much higher (possibly infinite) dimension via the mapping $\phi: \mathcal{X} \rightarrow \mathcal{F}$, and then projections are carried out in \mathcal{F} , commonly referred to as the feature space. The mapping ϕ is implicitly defined by the covariance kernel $k(\cdot, \cdot)$ as any symmetric, positive, real kernel can be expressed in terms of an inner product so that $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi_i, \phi_j \rangle$, where $\phi_i \triangleq \phi(\mathbf{x}_i)$ is the image of \mathbf{x}_i in the feature space [20]. For example, the feature space image of $x_i \in \mathbb{R}$ that corresponds to $k(x_i, x_j) = \exp\{-\frac{1}{2}(x_i - x_j)^2\}$ is an infinite dimensional vector whose m^{th} element is $[\phi_i]_m = \frac{x_i}{\sqrt{m!}} \exp\{-\frac{1}{2}x_i^2\}$. We remark that an explicit form for the mapping $\phi: \mathcal{X} \rightarrow \mathcal{F}$ is not needed because all algorithmic computations are expressed in terms of the inner product $\langle \phi_i, \phi_j \rangle = k(\mathbf{x}_i, \mathbf{x}_j)$.

With t data points $\{(\mathbf{x}_i, z_i)\}_{i=1}^t$ stored in memory, suppose a new data point $(\mathbf{x}_{t+1}, z_{t+1})$ has been collected, then it follows from the linearity of \mathcal{F} that ϕ_{t+1} can be written as the sum of two components

$$\phi_{t+1} = \underbrace{\sum_{i=1}^t \hat{\mathbf{e}}_{t+1}(i) \phi_i}_{\triangleq \hat{\phi}_{t+1}} + \gamma_{t+1} \phi_{res}.$$

namely, the orthogonal projection $\hat{\phi}_{t+1}$ and the residual, which is orthogonal to the previous t feature space images and lies along the unit vector ϕ_{res} . The expression for the linear coefficients

$$\hat{\mathbf{e}}_{t+1} = \mathbf{K}_t^{-1} \mathbf{k}_t(\mathbf{x}_{t+1}) \quad (5)$$

is derived so as to minimize the squared length of the residual $\gamma_{t+1} = \|\hat{\phi}_{t+1} - \phi_{t+1}\|_{\mathcal{F}}^2$, with the feature space norm given by $\|\phi(\mathbf{x})\|_{\mathcal{F}}^2 = \langle \phi(\mathbf{x}), \phi(\mathbf{x}) \rangle = k(\mathbf{x}, \mathbf{x})$. Thus, it has been shown [13] that γ_{t+1} simplifies to

$$\gamma_{t+1} = k(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}) - \mathbf{k}_t(\mathbf{x}_{t+1})^\top \mathbf{K}_t^{-1} \mathbf{k}_t(\mathbf{x}_{t+1}). \quad (6)$$

Upon comparison with (1b), γ_{t+1} can be interpreted as the predictive variance at \mathbf{x}_{t+1} given the noise-free observations at $\{\mathbf{x}_1, \dots, \mathbf{x}_t\}$.

Intuitively, if $\gamma_{t+1} = 0$, then ϕ_{t+1} can be expressed precisely as a linear combination of the previous t feature space images, and \mathbf{x}_{t+1} need not be stored in memory. In general, however, $\gamma_{t+1} \neq 0$. In fact, for the squared exponential covariance kernel (2), any set of distinct points maps to a linearly independent set of feature vectors [19]. Nonetheless, γ_{t+1} can be used as a heuristic to determine when a new data point should be explicitly retained in memory. If γ_{t+1} exceeds

a user-defined threshold ε_{tol} , then the new point $(\mathbf{x}_{t+1}, z_{t+1})$ is considered sufficiently novel to be stored in memory and the variables for prediction $\boldsymbol{\alpha}_t$ and \mathbf{C}_t undergo a standard update

$$\boldsymbol{\alpha}_{t+1} = \begin{bmatrix} \boldsymbol{\alpha}_t \\ 0 \end{bmatrix} + q_{t+1} \mathbf{s}_{t+1}, \quad (7a)$$

$$\mathbf{C}_{t+1} = \begin{bmatrix} \mathbf{C}_t & \mathbf{0}_t \\ \mathbf{0}_t^\top & 0 \end{bmatrix} + r_{t+1} \mathbf{s}_{t+1} \mathbf{s}_{t+1}^\top, \quad (7b)$$

$$\mathbf{s}_{t+1} = \begin{bmatrix} \mathbf{C}_t \mathbf{k}_t(\mathbf{x}_{t+1}) \\ 0 \end{bmatrix} + \mathbf{e}_{t+1}, \quad (7c)$$

where $\mathbf{e}_{t+1} = [\mathbf{0}_t^\top, 1]^\top$. However, if $\gamma_{t+1} \leq \varepsilon_{tol}$ then the data point is considered redundant and discarded. Prior to discarding the point, $\boldsymbol{\alpha}_t$ and \mathbf{C}_t undergo a reduced update, similar to the standard update, but the dimensionality of the variables is not extended, and \mathbf{e}_{t+1} is replaced by the orthogonal projection coefficients $\hat{\mathbf{e}}_{t+1}$

$$\boldsymbol{\alpha}_{t,upd} = \boldsymbol{\alpha}_t + q_{t+1} \hat{\mathbf{s}}_{t+1}, \quad (8a)$$

$$\mathbf{C}_{t,upd} = \mathbf{C}_t + r_{t+1} \hat{\mathbf{s}}_{t+1} \hat{\mathbf{s}}_{t+1}^\top, \quad (8b)$$

$$\hat{\mathbf{s}}_{t+1} = \mathbf{C}_t \mathbf{k}_t(\mathbf{x}_{t+1}) + \hat{\mathbf{e}}_{t+1}, \quad (8c)$$

with the scalars r_{t+1} and q_{t+1} given by

$$q_{t+1} = \frac{z_{t+1} - \mathbf{k}_t(\mathbf{x}_{t+1})^\top \boldsymbol{\alpha}_t}{\sigma_n^2 + k(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}) + \mathbf{k}_t(\mathbf{x}_{t+1})^\top \mathbf{C}_t \mathbf{k}_t(\mathbf{x}_{t+1})},$$

$$r_{t+1} = \frac{-1}{\sigma_n^2 + k(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}) + \mathbf{k}_t(\mathbf{x}_{t+1})^\top \mathbf{C}_t \mathbf{k}_t(\mathbf{x}_{t+1})}.$$

Even though the computation of $\hat{\mathbf{e}}_{t+1}$ (5) and γ_{t+1} (6) requires the inverse of the gram matrix \mathbf{K}_t , the costly inversion is avoided by recursively updating $\mathbf{Q}_t \triangleq \mathbf{K}_t^{-1}$ via

$$\mathbf{Q}_{t+1} = \begin{bmatrix} \mathbf{Q}_t & \mathbf{0}_t \\ \mathbf{0}_t^\top & 0 \end{bmatrix} + \gamma_{t+1}^{-1} \begin{bmatrix} \hat{\mathbf{e}}_{t+1} \\ -1 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{e}}_{t+1}^\top \\ -1 \end{bmatrix}^\top, \quad (9)$$

whenever a new data point is added to the dataset.

Based on available computing resources, the user can specify an upper limit N_{max} on the number of data points retained for the regression model. As soon as the maximum number of data points has been collected, a decision rule is needed to determine if a new sufficiently novel measurement $(\mathbf{x}_{t+1}, z_{t+1})$ should replace a measurement already stored in the dataset. To do so, the algorithm first assumes the new data point has been added to the dataset, and then each point in the dataset is assigned a score based on the change of the posterior mean due to projection. As points that induce a smaller change are seen as less informative, the point with the smallest score is removed from the dataset.

For an input $\mathbf{x} \in \mathcal{X}$, let $\mu_{t+1}(\mathbf{x})$ denote the mean value that arises due to the standard update (7) with $(\mathbf{x}_{t+1}, z_{t+1})$ and let $\hat{\mu}_{t+1}(\mathbf{x})$ denote the mean value that arises due to the reduced update (8), then the change in posterior mean due to projection evaluated at \mathbf{x} is $\Delta\mu_{t+1}(\mathbf{x}) = |\mu_{t+1}(\mathbf{x}) - \hat{\mu}_{t+1}(\mathbf{x})|$. The score for $(\mathbf{x}_{t+1}, z_{t+1})$ is the sum of the changes in posterior mean evaluated at all the locations $\{\mathbf{x}_1, \dots, \mathbf{x}_{t+1}\}$, specifically $\varepsilon_{t+1} = \sum_{i=1}^{t+1} \Delta\mu_{t+1}(\mathbf{x}_i)$, which has been shown

[13] to reduce to

$$\varepsilon_{t+1} = |q_{t+1}| \gamma_{t+1} = \frac{|[\boldsymbol{\alpha}_{t+1}]_{t+1}|}{|[\mathbf{Q}_{t+1}]_{t+1,t+1}|}.$$

Under the assumption that the order of the data has no influence on the posterior predictive distribution, the score for the i^{th} data point is

$$\varepsilon_i = \frac{|[\boldsymbol{\alpha}_{t+1}]_i|}{|[\mathbf{Q}_{t+1}]_{i,i}|}. \quad (10)$$

The cost of computing the scores is $\mathcal{O}(t+1)$, rendering the scores computationally inexpensive.

Let (\mathbf{x}_r, z_r) denote the data point with the lowest score. To remove (\mathbf{x}_r, z_r) from the dataset, first reorder the elements of $\boldsymbol{\alpha}_{t+1}, \mathbf{C}_{t+1}, \mathbf{Q}_{t+1}$ so that they are consistent with (\mathbf{x}_r, z_r) being the last element in the dataset, and then define

$$\begin{aligned} \boldsymbol{\alpha}_{t+1} &= \begin{bmatrix} [\boldsymbol{\alpha}_{t+1}]_{1:t} \\ \alpha_* \end{bmatrix}, \\ \mathbf{C}_{t+1} &= \begin{bmatrix} [\mathbf{C}_{t+1}]_{1:t,1:t} & \mathbf{c}_{t+1} \\ \mathbf{c}_{t+1}^\top & c_* \end{bmatrix}, \\ \mathbf{Q}_{t+1} &= \begin{bmatrix} [\mathbf{Q}_{t+1}]_{1:t,1:t} & \mathbf{q}_{t+1} \\ \mathbf{q}_{t+1}^\top & q_* \end{bmatrix}, \end{aligned}$$

so as to compute

$$\hat{\boldsymbol{\alpha}}_t = [\boldsymbol{\alpha}_{t+1}]_{1:t} - \frac{\alpha_*}{q_*} \mathbf{q}_{t+1}, \quad (11a)$$

$$\begin{aligned} \hat{\mathbf{C}}_t &= [\mathbf{C}_{t+1}]_{1:t,1:t} + \frac{c_*}{q_*^2} \mathbf{q}_{t+1} \mathbf{q}_{t+1}^\top \\ &\quad - \frac{1}{q_*} [\mathbf{q}_{t+1} \mathbf{c}_{t+1}^\top + \mathbf{c}_{t+1} \mathbf{q}_{t+1}^\top] \end{aligned} \quad (11b)$$

$$\hat{\mathbf{Q}}_t = [\mathbf{Q}_{t+1}]_{1:t,1:t} - \frac{1}{q_*} \mathbf{q}_{t+1} \mathbf{q}_{t+1}^\top, \quad (11c)$$

where $\hat{\boldsymbol{\alpha}}_t, \hat{\mathbf{C}}_t, \hat{\mathbf{Q}}_t$ are the variables with (\mathbf{x}_r, z_r) removed from the dataset.

The update procedure for the prediction variables used in sparse online Gaussian process regression is summarized in Algorithm 1.

IV. MULTI-AGENT COMMUNICATION APPROACH

For the ideal scenario in which agents can communicate instantaneously, we propose the straightforward approach that an agent broadcasts to all other agents whenever the agent gathers a new measurement that has been retained in the data set, meaning that (i) the measurement is sufficiently novel $\gamma_{t+1} > \varepsilon_{tol}$, and (ii) the measurement is not immediately removed (i.e. $\varepsilon_{t+1} \neq \min\{\varepsilon_1, \dots, \varepsilon_{t+1}\}$ when $t+1$ exceeds N_{max}).

For the more practical scenario in which communication is range-limited, agents must enter into communication range to share the most recent additions to their dataset. We consider the case in which an agent's most recent set of retained measurements is too large to be transmitted in its entirety, and so we propose the approach that the agent transmit an informative subset of their novel measurements.

Let $S_t = \{(\mathbf{x}_i, z_i)\}_{i=1}^t$ denote the measurements that have yet to be communicated by an agent, and suppose the communication bandwidth dictates that only $k < t$ measurements

input : Novelty threshold ε_{tol} , maximum number of data points N_{max}

output: Regression variables $\mathbf{X}, \mathbf{z}, \boldsymbol{\alpha}, \mathbf{C}, \mathbf{Q}$

```

1  $\mathbf{X}_t \leftarrow \emptyset, \mathbf{z}_t \leftarrow \emptyset, \boldsymbol{\alpha}_t \leftarrow \emptyset, \mathbf{C}_t \leftarrow \emptyset, \mathbf{Q}_t \leftarrow \emptyset;$ 
2 while collecting measurements do
3    $(\mathbf{x}_{t+1}, z_{t+1}) \leftarrow$  Obtain measurement;
4    $\gamma_{t+1} \leftarrow$  Compute novelty metric (6);
5   if  $\gamma_{t+1} < \varepsilon_{tol}$  then
6      $[\boldsymbol{\alpha}_{t,upd}, \mathbf{C}_{t,upd}] \leftarrow$  Reduced update (8);
7   else
8      $[\boldsymbol{\alpha}_{t+1}, \mathbf{C}_{t+1}, \mathbf{Q}_{t+1}] \leftarrow$  Std. update (7), (9);
9      $S_{t+1} \leftarrow S_t \cup (\mathbf{x}_{t+1}, z_{t+1});$ 
10    if  $t+1 > N_{max}$  then
11      for  $x_i \in \mathbf{X}_{t+1}$  do
12         $\varepsilon_i \leftarrow$  compute score (10);
13      end
14       $(\mathbf{x}_r, z_r) \leftarrow \arg \min_{(x_i, z_i)} \varepsilon_i;$ 
15       $[\hat{\boldsymbol{\alpha}}_t, \hat{\mathbf{C}}_t, \hat{\mathbf{Q}}_t] \leftarrow$  Remove  $(\mathbf{x}_r, z_r)$  (11);
16       $S_t \leftarrow S_{t+1} \setminus (\mathbf{x}_r, z_r);$ 
17    end
18  end
19 end

```

Algorithm 1: The update procedure for the prediction variables used in sparse online Gaussian process regression.

can be communicated. The task is to select a subset $S_k \subset S_t$ with the intention of optimizing a user-selected criterion. For criteria we consider the conventional information-theoretic measures entropy and mutual information, and we consider a new measurement-based criterion, namely the score (10) derived from the sparse online Gaussian process algorithm.

Criterion 1 (Entropy): The uncertainty of a random variable is formalized by the concept of entropy [21, Ch. 9.1]. Here, the aim is to select the subset $S_k \subset S_t$ that minimizes the uncertainty of the predictions at the remaining unshared locations, given that S_k has been shared, characterized by the conditional entropy $H(S_t \setminus S_k | S_k)$. From the chain rule of conditional entropy, it follows that $H(S_t \setminus S_k | S_k) = H(S_t) - H(S_k)$, which means that minimizing the conditional entropy is equivalent to maximizing the entropy $H(S_k)$. Unfortunately, the optimization problem

$$S_k = \arg \max_{S'_k \subset S_t : |S'_k| \leq k} H(S'_k)$$

has been proven [22] to be NP-hard. Thus, a greedy selection algorithm is commonly employed [2]. The algorithm begins with an empty set $S_0 = \emptyset$ and at the j^{th} iteration adds the data point $s = (\mathbf{x}_s, z_s) \in S_t \setminus S_{j-1}$ that has the highest conditional entropy

$$H(s|S_{j-1}) = \frac{1}{2} \ln(2\pi e \sigma_{\mathbf{x}_s|S_{j-1}}^2), \quad (12)$$

where $\sigma_{\mathbf{x}_s|S_{j-1}}^2$ is computed according to (1b). The greedy selection algorithm is summarized in Algorithm 2.

input : $S_t = (\mathbf{X}_t, \mathbf{z}_t)$, subset size k
output: Subset $S_k \subseteq S_t$

```

1  $S_0 \leftarrow \emptyset, \mathbf{X} \leftarrow \mathbf{X}_t;$ 
2 for  $j = 1 : k$  do
3   for  $x_i \in \mathbf{X}$  do
4      $H_{s_i} \leftarrow$  compute conditional entropy (12);
5   end
6    $(\mathbf{x}, z) \leftarrow \arg \max_{(x_i, z_i)} H_{s_i};$ 
7    $S_j \leftarrow S_{j-1} \cup (\mathbf{x}, z);$ 
8    $\mathbf{X} \leftarrow \mathbf{X} \setminus \mathbf{x};$ 
9 end

```

Algorithm 2: Algorithm to greedily select the most informative subset of data points based on entropy.

Criterion 2 (Mutual Information): Alternatively, as initially proposed in [23], S_k can be selected so as to maximize the mutual information $MI(S_k) = H(S_t \setminus S_k) - H(S_t \setminus S_k | S_k)$ between the locations contained in S_k and the locations contained in the remainder $S_t \setminus S_k$. As was the case for entropy, the optimal solution cannot, in general, be computed in polynomial time and so a greedy algorithm is employed to obtain an approximation of the optimal solution [2]. The algorithm begins with an empty set $S_0 = \emptyset$ and at the j^{th} iteration adds the data point $s \in S_t \setminus S_{j-1}$ that yields the largest increase in mutual information

$$\underbrace{MI(S_{j-1} \cup s) - MI(S_{j-1})}_{\triangleq \delta_s} = H(s | S_{j-1}) - H(s | \bar{S}_{j-1}), \quad (13)$$

where $\bar{S}_{j-1} \triangleq S_t \setminus (S_{j-1} \cup s)$ has been defined for notational convenience, and the requisite entropies are computed with (12). The algorithm to determine S_k is given in Algorithm 3.

input : $S_t = (\mathbf{X}_t, \mathbf{z}_t)$, subset size k
output: Subset $S_k \subseteq S_t$

```

1  $S_0 \leftarrow \emptyset, \mathbf{X} \leftarrow \mathbf{X}_t;$ 
2 for  $j = 1 : k$  do
3   for  $x_i \in \mathbf{X}$  do
4      $\delta_{s_i} \leftarrow$  compute MI gain (13);
5   end
6    $(\mathbf{x}, z) \leftarrow \arg \max_{(x_i, z_i)} \delta_{s_i};$ 
7    $S_j \leftarrow S_{j-1} \cup (\mathbf{x}, z);$ 
8    $\mathbf{X} \leftarrow \mathbf{X} \setminus \mathbf{x};$ 
9 end

```

Algorithm 3: Algorithm to greedily select the most informative subset of data points based on mutual information.

Criterion 3 (SOGP Score): Given that entropy and mutual information are a function of the conditional entropy (12), these measures reflect only the location at which a measurement was taken and not the value of the measurement itself. Thus, we propose a new criterion which reflects the value of the measurement.

A measurement that agrees poorly with the predictive mean (4a) will give rise to substantial change in predictive model and is thus associated with a high gain in information. The score (10), derived from the sparse online Gaussian process regression algorithm, rigorously quantifies the information gain of each data point. The algorithm to select the most informative subset $S_k \subset S_t$ begins with the entire set of measurements S_t and iteratively removes the point (\mathbf{x}_r, z_r) with the smallest score until there are only k remaining data points. The algorithmic details are listed in Algorithm 4.

input : Regression variables
 $S_t = (\mathbf{X}_t, \mathbf{z}_t), \boldsymbol{\alpha}_t, \mathbf{C}_t, \mathbf{Q}_t$, subset size k
output: Subset $S_k \subseteq S_t$

```

1 for  $j = 1 : t - k$  do
2   for  $x_i \in \mathbf{X}_{t-j+1}$  do
3      $\varepsilon_i \leftarrow$  compute score (10);
4   end
5    $(\mathbf{x}_r, z_r) \leftarrow \arg \min_{(x_i, z_i)} \varepsilon_i;$ 
6    $[\hat{\boldsymbol{\alpha}}_{t-j}, \hat{\mathbf{C}}_{t-j}, \hat{\mathbf{Q}}_{t-j}] \leftarrow$  Remove  $(\mathbf{x}_r, z_r)$  (11);
7    $S_{t-j} \leftarrow S_{t-j+1} \setminus (\mathbf{x}_r, z_r);$ 
8 end

```

Algorithm 4: Algorithm to select the most informative subset of data points based on the sparse online Gaussian process score (10).

V. EMPIRICAL ASSESSMENT

To evaluate our approach to reduce communication, we consider the scenario in which a single agent collects all the measurements in a given dataset and must judiciously select which measurements to share with the other agents. Special emphasis is placed on the number of shared measurements with respect to the total number of measurements. Thus, all of the datasets used to empirically assess our multi-agent communication approach contain thousands of measurements. We selected a collection of real-world data sets that are relevant to robotic mapping applications. The collection consists of a bathymetry dataset, a magnetic field dataset, and a radio signal strength (RSS) dataset.

The bathymetry dataset was collected at Claytor Lake, in Dublin, Virginia, USA, located near the Virginia Tech Blacksburg campus, using a depth-sounder on board a GPS-equipped boat. The dataset spans an area of 2000m x 2000m and contains over 13,000 measurements that take on values in the range of 1 to 35m in depth. The magnetic field dataset is known as the UJIIndoorLoc-Mag database [24]. This publicly available dataset was established as a benchmark dataset to compare the performance of magnetic field-based indoor localization methods. Measurements were taken along paths that cover approximately 650 lineal meters of an office, and the magnitude of the magnetic field measurements ranges from 3.82 to 77.44 μT . Lastly, the radio signal strength dataset consists of measurements collected by a mobile robot as it interacted with a stationary wireless access point. The

dataset is publicly available on the CRAWDAD repository, hosted by Dartmouth [25].

For the purposes of evaluating predictive performance, each dataset is randomly partitioned into a training set (80%) to sequentially train the regression model and a test set (20%) to assess predictive performance quantified in terms of root mean squared error (RMSE).

We applied the sparse online Gaussian process regression algorithm to each dataset using the hyperparameters listed in Table I. In our experience, we find that $\epsilon_{tol} < 0.00064\sigma_f$ leads to an ill-conditioned gram matrix \mathbf{K}_t , so we put $\epsilon_{tol} = 0.0064\sigma_f$ in our experiments, and chose N_{max} large enough so that no novel measurements need to be discarded, i.e. $N_{max} = 500$. Each dataset corresponds to a particular plot in Figure 1 wherein the locations of all measurements are depicted using light blue markers and sufficiently novel measurements are depicted using dark blue markers. Each plot lists the total number of measurements in the dataset along with the number of sufficiently novel measurements retained for regression. In all cases, the number of retained measurements is two orders of magnitude less than the total number of measurements. Moreover, regression with the retained measurements leads to a root mean squared error of 2.02m, 5.95 μ T, and 2.40 dBm for the bathymetric, magnetic field intensity, and radio signal strength dataset, respectively. Comparison of the root mean squared error with the measurement noise σ_n values in Table I suggests that the improvement to predictive performance by sharing additional measurements would be negligible.

TABLE I
REGRESSION HYPERPARAMETERS FOR EACH DATASET

Dataset	ℓ	σ_f	σ_n
Bathymetry	154.97	9.00	2.22
Magnetic intensity	2.21	13.76	6.37
Radio signal strength	9.56	15.19	2.58

In the case of range-limited communication, we consider the case that a single agent has collected all the measurements in a dataset without being in communication range of any other agents and has performed sparse Gaussian process regression to retain only those points that are sufficiently novel and must resort to sharing only a subset of the sufficiently novel measurements. Here, we empirically evaluate the subset selection algorithms based on entropy, mutual information and the sparse online Gaussian process (SOGP) score.

The results in Figure 2 demonstrate that subset selection based the SOGP score leads to predictive performance that is no worse than subset selection based on the conventional information-theoretic metrics, entropy and mutual information. In fact, subset selection based on the SOGP score leads to superior predictive performance in most instances, especially for subsets of smaller size. The score (10) reflects not only the spatial correlation of a measurement with respect to all other measurements, but also reflects how informative a measurement is, i.e. the difference between the measurement

and the predictive mean. Given that entropy and mutual information do not leverage the information encoded in measurements, it is perhaps consistent with intuition that subset selection based on the SOGP score yields superior predictive performance.

VI. CONCLUSION

Our empirical findings demonstrate that our approach to communication enables an agent to substantially reduce, in a rigorous manner, the amount of data it shares with other agents. However, the simulation studies are in the context of a single agent making a decision as to which measurements to communicate. For the findings to be more broadly applicable, it is of interest to incorporate our communication approach into a coordination strategy, such as those proposed in [4]–[8], and perform studies to assess whether our approach provides timely information to other agents that is relevant to their tasks.

A key limitation of our approach is that the novelty measure relies on a fixed set of hyperparameters $\theta = \{\ell, \sigma_f, \sigma_n\}$. Loosely speaking, ℓ characterizes the amount of distance in input space that must be traveled before the function value can change significantly. To address cases in which regions of the spatial field are best characterized by different length-scale parameters or different measurement noise levels, it would be ideal to have a rigorous novelty metric that can cope with sets of hyperparameters for various regions of the input domain or sets of hyperparameters that evolve in space. Then the values of the hyperparameters along with the novelty metric could inform the coordination of other agents. For example, if an agent discovers an area characterized by a highly-varying spatial field with a high noise level, then it could alert other agents to assist in mapping that area. This is in contrast to an area characterized by a constant spatial field with a low noise level which would need fewer measurements to adequately map.

Additionally, our empirical results show that the measurement-based criterion (10) for subset selection leads to improved predictive performance. As the score reflects the prediction error, it may be of value to explore its use in informing path-planning and coordination. Ideally, this would drive agents to collect measurements in regions of the input space that are poorly modeled so as to best improve the prediction model.

REFERENCES

- [1] S. Vasudevan, F. Ramos, E. Nettleton, and H. Durrant-Whyte, "Gaussian process modeling of large-scale terrain," *Journal of Field Robotics*, vol. 26, no. 10, pp. 812–840, 2009.
- [2] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," *Journal of Machine Learning Research*, vol. 9, no. Feb, pp. 235–284, 2008.
- [3] B. Ferris, D. Fox, and N. D. Lawrence, "Wifi-slam using gaussian process latent variable models," in *IJCAI*, vol. 7, no. 1, 2007, pp. 2480–2485.
- [4] A. Viseras, Z. Xu, and L. Merino, "Distributed multi-robot information gathering under spatio-temporal inter-robot constraints," *Sensors*, vol. 20, no. 2, p. 484, 2020.

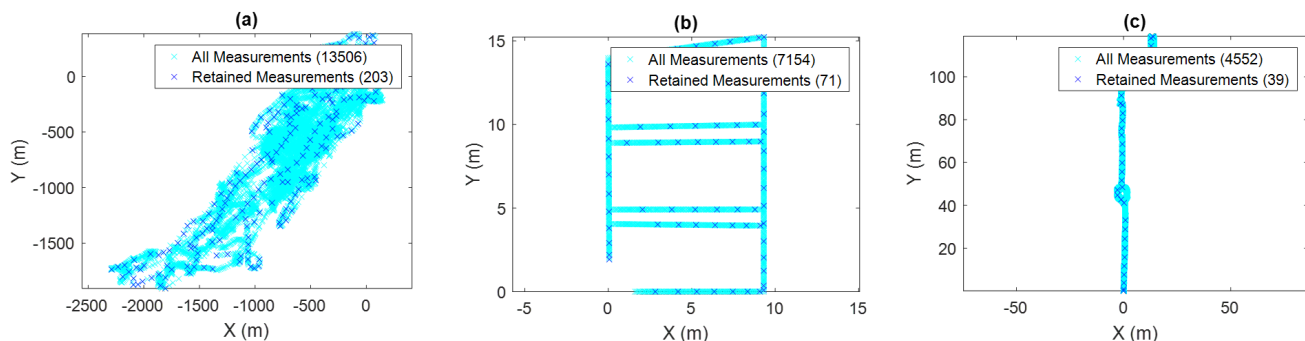


Fig. 1. Measurement locations pertaining to the bathymetry dataset (a), magnetic field dataset (b), and radio signal strength dataset (c). The light blue markers denote all measurements in the dataset and the dark blue markers denote the measurements that have been retained for regression in accordance with the sparse online Gaussian process regression algorithm. These empirical results demonstrate that an agent can significantly reduce the number of shared measurements by broadcasting only those measurements that are retained for regression.

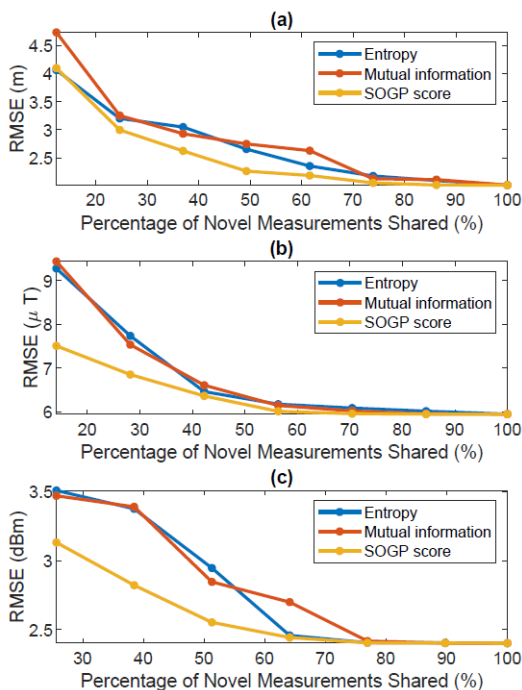


Fig. 2. Predictive performance in the context of an agent selecting the most informative subset of data to broadcast to other agents based on entropy, mutual information, and the sparse online Gaussian process SOGP score (10), evaluated on the bathymetric dataset (a), magnetic field dataset (b), and radio signal strength dataset (c). In the majority of instances, the SOGP score leads to superior predictive performance.

[5] A. Viseras, T. Wiedemann, C. Manss, L. Magel, J. Mueller, D. Shutin, and L. Merino, "Decentralized multi-agent exploration with online-learning of gaussian processes," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4222–4229.

[6] D.-H. Cho, J.-S. Ha, S. Lee, S. Moon, and H.-L. Choi, "Informative path planning and mapping with multiple uavs in wind fields," in *Distributed Autonomous Robotic Systems*. Springer, 2018, pp. 269–283.

[7] J. Banfi, A. Q. Li, N. Basilico, I. Rekleitis, and F. Amigoni, "Multi-robot online construction of communication maps," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2577–2583.

[8] R. Ouyang, K. H. Low, J. Chen, and P. Jaillet, "Multi-robot active sensing of non-stationary gaussian process-based environmental phe-

nomena," 2014.

[9] A. Viseras, Z. Xu, and L. Merino, "Distributed multi-robot cooperation for information gathering under communication constraints," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1267–1272.

[10] M. N. Rooker and A. Birk, "Multi-robot exploration under the constraints of wireless networking," *Control Engineering Practice*, vol. 15, no. 4, pp. 435–445, 2007.

[11] R. Pandey, A. K. Singh, and K. M. Krishna, "Multi-robot exploration with communication requirement to a moving base station," in *2012 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2012, pp. 823–828.

[12] G. Han, C. Zhang, L. Shu, and J. J. Rodrigues, "Impacts of deployment strategies on localization performance in underwater acoustic sensor networks," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 3, pp. 1725–1733, 2014.

[13] L. Csato and M. Opper, "Sparse on-line gaussian processes," *Neural computation*, vol. 14, no. 3, pp. 641–668, 2002.

[14] O. Amir, B. J. Grosz, R. Stern, *et al.*, "To share or not to share? the single agent in a team decision problem." 2014.

[15] M. Roth, R. Simmons, and M. Veloso, "What to communicate? execution-time decision in multi-agent pomdps," in *Distributed Autonomous Robotic Systems 7*. Springer, 2006, pp. 177–186.

[16] F. Wu, S. Zilberstein, and X. Chen, "Online planning for multi-agent systems with bounded communication," *Artificial Intelligence*, vol. 175, no. 2, pp. 487–511, 2011.

[17] G. Best, M. Forrai, R. R. Mettu, and R. Fitch, "Planning-aware communication for decentralised multi-robot coordination," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1050–1057.

[18] R. J. Marcotte, X. Wang, D. Mehta, and E. Olson, "Optimizing multi-robot communication under bandwidth constraints," *Autonomous Robots*, vol. 44, no. 1, pp. 43–55, 2020.

[19] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[20] B. Scholkopf, S. Mika, C. J. Burges, P. Knirsch, K.-R. Muller, G. Ratsch, and A. J. Smola, "Input space versus feature space in kernel-based methods," *IEEE transactions on neural networks*, vol. 10, no. 5, pp. 1000–1017, 1999.

[21] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 1991.

[22] C.-W. Ko, J. Lee, and M. Queyranne, "An exact algorithm for maximum entropy sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, 1995.

[23] W. F. Caselton and J. V. Zidek, "Optimal monitoring network designs," *Statistics & Probability Letters*, vol. 2, no. 4, pp. 223–227, 1984.

[24] J. Torres-Sospedra, D. Rambla, R. Montoliu, O. Belmonte, and J. Huerta, "Ujiindoorloc-mag: A new database for magnetic field-based localization problems," in *2015 International conference on indoor positioning and indoor navigation (IPIN)*. IEEE, 2015, pp. 1–10.

[25] R. Parasuraman, S. Caccamo, F. Baberg, and P. Ogren, "CRAWDAD dataset kth/rss (v. 2016-01-05)," Downloaded from <https://crawdad.org/kth/rss/20160105/outdoor>, Jan. 2016, traceset: outdoor.